

Going Deeper for Multilingual Visual Sentiment Detection

Brendan Jou and Shih-Fu Chang

Department of Electrical Engineering, Columbia University, New York, NY 10027

Keywords: Sentiment; Ontology; Multilingual; Concept Detection; Deep Networks; Fine-tuning

ABSTRACT

This technical report details several improvements to the visual concept detector banks built on images from the Multilingual Visual Sentiment Ontology (MVSO).¹ The detector banks are trained to detect a total of 9,918 sentiment-biased visual concepts from six major languages: English, Spanish, Italian, French, German and Chinese. In the original MVSO release,¹ adjective-noun pair (ANP) detectors were trained for the six languages using an AlexNet-styled architecture² by fine-tuning from DeepSentiBank.³ Here, through a more extensive set of experiments, parameter tuning, and training runs, we detail and release higher accuracy models for detecting ANPs across six languages from the same image pool and setting as in the original release using a more modern architecture, GoogLeNet,⁴ providing comparable or better performance with reduced network parameter cost.

In addition, since the image pool in MVSO¹ can be corrupted by user noise from social interactions, we partitioned out a sub-corpus of MVSO images based on tag-restricted queries for higher fidelity labels. We show that as a result of these higher fidelity labels, higher performing AlexNet-styled² ANP detectors can be trained using the tag-restricted image subset as compared to the models in full corpus. We release all these newly trained models for public research use along with the list of tag-restricted images from the MVSO dataset.

1. INTRODUCTION

Following the trend in many other fields, the advent of high-volume and weakly-supervised data is driving increased interest in *large-scale* sentiment studies in Affective Computing.⁵ However, directly studying affect, or emotion, in a dimensional representation (e.g. valence-arousal-dominance) or even discrete semantics (e.g. sad, happy, angry, etc.) tend to suffer from the problem of data sparsity since such specialized psychology terminology are less likely to be found in large volume from the Web. Inspired by similar work in Computer Vision and Multimedia, several works have since proposed affective mid-level representations to bridge the *affective gap* between low-level features and high-level affect states. One such attractive mid-level semantic representation is the *adjective-noun pair* (ANP),⁶ where ‘nouns’ are used to provide visual grounding to a detectable object and ‘adjectives’ are used to sentimentally bias the object.

In a recent work, we presented a large-scale multilingual visual sentiment ontology (MVSO)¹ consisting of 15,630 such ANPs across 12 languages with an associated dataset of over 7.36 million images. As part of this initial work, we developed a set of language-specific visual concept detector banks that could identify the presence of ANPs in six of the major languages in the MVSO. These detector banks were shown to be useful in cross-lingual sentiment analysis,¹ and recently, in diversifying image query expansion⁷ and facial imagery understanding.⁸ Yet in the original development of the ANP detectors, many of the original training parameter settings were defaulted and left as-is for a proof-of-concept. Even though using such default settings show that the multilingual ANP detectors can achieve reasonable detection performance already, they do not represent the most near-optimal networks that can be trained for the given architectures used.

In this report, we detail and release (1) higher accuracy models for detecting adjective-noun pairs (ANPs) in six languages from the same image pool in the original Multilingual Visual Sentiment Ontology (MVSO) release¹ and new detector banks fine-tuned using a different more modern network architecture with benchmark comparisons, (2) a sub-corpus of MVSO based on tag-restricted queries for higher fidelity labels, and (3) adjective-noun pair detectors based on this tag-restricted subset. The model and data release can be found through the original MVSO website at <http://mvso.cs.columbia.edu>.

Primary correspondence author: B. Jou.

E-mails: {bjou,sfchang}@ee.columbia.edu

2. BRIEF OVERVIEW OF MVSO

The Multilingual Visual Sentiment Ontology (MVSO),¹ a substantial extension and improvement on the Visual Sentiment Ontology (VSO),⁶ consists of over 15.6K sentiment-biased mid-level visual concepts using the semantic construct of adjective-noun pairs (ANPs). The ontology was constructed by first using seed emotion keywords from “Plutchik’s Wheel of Emotions”⁹ where were translated into 12 different languages by native speakers. For each language, these keywords were used to query the Flickr API* to retrieve a large corpus of images with related tags and other metadata. From the tags and metadata, we mined popularly used adjectives and nouns, identified by language-dependent part-of-speech taggers, combinatorially formed adjective-noun pair (ANP) candidates, and then filtered them on a number of criteria including semantic correctness, sentiment strength, popular usage on Flickr, and coverage in Flickr uploader diversity. After this candidate filtering, the MVSO consisted of 15,630 ANP concepts representative of 12 different languages (language codes in brackets): Arabic [ar], Chinese [zh], Dutch [nl], English [en], French [fr], German [de], Italian [it], Persian [fa], Polish [pl], Russian [ru], Spanish [es] and Turkish [tr]. These ANP concepts were then used to query the Flickr API to retrieve associated images and metadata, and it was later found that our ANPs also have geo-reference coverage of over 235 countries.¹⁰

To develop an actual ontology with a semantic tree structure, we proposed two construction methods: *exact matching* and *approximate matching*.^{1,8} In ‘exact matching,’ we use a pivot language where all ANPs are automatically translated into English and ANPs become naturally grouped by their translations. This yields high cluster “precision,” but low “recall” since nuances like plurals and synonyms, and even related concepts are not grouped together. In ‘approximate matching,’ we extract word feature vectors from the ANPs and perform traditional distance-based clustering. This trades of some cluster consistency for increased cluster coverage. These clustering approaches have proven useful in uncovering some multilingual geographic social phenomena¹⁰ as well as in cross-lingual sentiment analysis and query expansion.^{1,7} Samples images from MVSO for an ANP cluster with concepts like ‘old house’ are shown in Figure 1.

The final MVSO image corpus consists of 7,368,364 images from which language-specific ANP detector banks were trained. Six of the major languages were chosen to train these detectors, resulting in a coverage of 9,918 total ANPs across 6,994,319 images in English [en], Spanish [es], Italian [it], French [fr], German [de] and Chinese [zh]. For more details on MVSO, we refer the interested reader to the related peer-review publications.^{1,7,8,10}

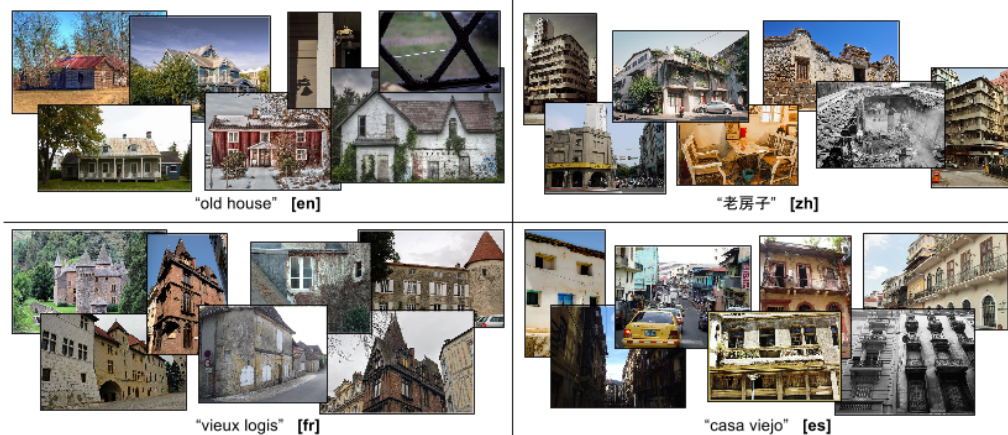


Figure 1: Example MVSO Images in Four Sampled Languages from a Multilingual ANP Cluster. The cluster, formed by approximate matching, portrays concepts related to the “old” or “ancient” nature of some “houses”.

3. DEEP VISUAL SENTIMENT CONCEPT DETECTOR BANKS

In our original MVSO publication,¹ we trained six detector banks, fine-tuning using an AlexNet-styled network architecture where there are five convolutional layers and three fully-connected layers² with pooling and normalization layers swapped to create a variant called CaffeNet.¹¹ We initialized our networks weights from

*<https://www.flickr.com/services/api>

Language	#ANPs	#params	#train	#test	time	top-1	top-5
English	4,342	74.66	3,236,728	807,447	95	10.13%	21.06%
Spanish	2,382	66.63	1,085,678	270,400	45	12.35%	25.36%
Italian	1,561	63.26	602,424	149,901	30	17.01%	30.93%
French	1,115	61.44	462,522	115,112	26	17.66%	35.46%
German	275	57.99	108,744	27,048	12	30.11%	52.78%
Chinese	243	57.86	102,740	25,575	15	27.07%	45.06%

Table 1: Classification accuracies of the ANP detector banks (%) from the original MVSO release,¹ where images were collected from a hybrid pool of tag-restricted and free text searches. This table essentially mirrors Table 4 in the original publication in ACM Multimedia 2015¹ except now to two-decimal precision and with estimated walltimes (hours) instead of CPU/GPU process timing and number of network parameters (millions).

a CaffeNet fine-tuned network called DeepSentiBank,³ which was trained to detect 2,089 ANPs (in English) from the images of the VSO dataset⁶ and improving on its original proposed SVM-based detector banks, SentiBank.¹² In fine-tuning these six detector banks, we followed the traditional fine-tuning strategy of resizing the final fully-connected layer to have the same number of output units as the number of classes with random initialization (in this case, Gaussian), started optimization at a lower base learning rate (in this case, 0.001), and used a higher learning rate multiplier for the last fully-connected layer to allow for more aggressive updates on that layer compared to the other layers that were pretrained (in this case, $\times 2$). During optimization, we used a batch size of 256 and fixed the learning rate, gradually reducing it by a factor of 10 whenever we saw training loss plateau. Every other setting in the network architecture, e.g. dropout rates and output sizes of other fully-connected layers, and training routines, e.g. optimization technique, was kept the same as when training the original ILSVRC¹³ model[†]. We note that in these MVSO detectors, the original ILSVRC12 mean image was used for mean subtraction. All training was originally performed on a single NVIDIA GeForce GTX 980 GPU and was implemented with Caffe.¹¹ The performances in the original MVSO release¹ are copied in Table 1.

To further improve the detection performance of ANPs, we experimented with a much deeper and more complex architecture called GoogLeNet, or Inception.⁴ The original GoogLeNet design consists of mini-networks, also called Inception modules, that are composed of $1\times 1C$, $1\times 1C-3\times 3C$, $1\times 1C-5\times 5C$, and $3\times 3MP-1\times 1C$ towers where C corresponds to convolutions and MP denotes max pooling, i.e. so $1\times 1C-3\times 3C$ denotes one path in the modules that has a 1×1 convolution followed by a 3×3 convolution. While AlexNet-styled architectures are only 8 layers deep, the GoogLeNet architecture is 22 layers deep while still using less parameters[‡]. These fan-in-fan-out modules are stacked in an architecture referred to as network-in-network.¹⁴ GoogLeNet additionally uses two auxiliary branches in the network to prevent gradients from vanishing during backpropagation.⁴

For training, we preserved the same training setting as in the original MVSO release¹ using the same train/test split of images. Unlike before though we pre-train from a ILSVRC12 model since the architecture is significantly different from DeepSentiBank³ and no comparable Inception-like network was previously trained on VSO. In addition, since the output space for languages like English and Spanish is large, in the auxiliary heads of Inception, we widened the second-to-last fully-connected layers which were originally 1,024 to 4,352 and 4,096, respectively, since they would otherwise become a representational bottleneck. For English and Spanish, we found that using the second auxiliary head as the final output actually yielded higher performance due to this layer widening and so we truncate the network after this second auxiliary head for the final network. For all other languages, the auxiliary heads were truncated and the usual penultimate layers were preserved in the final network. We used SGD with a sigmoid decay learning rate decay with a base learning rate of 0.001, a batch size of 64 and a decay factor of 0.1 and unlike before use the true training data mean image during mean subtraction.

4. EXPERIMENTS & DISCUSSION

All new experiments used a single NVIDIA GeForce GTX Titan X GPU and were implemented with Caffe.¹¹

[†]See https://github.com/BVLC/caffe/tree/master/models/bvlc_reference_caffenet for parameter details.

[‡]Note that here we use the first iteration of GoogLeNet/Inception, sometimes referred to as Inception-v1, though there have been several proposed improvements since that use double 3×3 layers or that integrate residual learning.

4.1 Hybrid-pool Multilingual ANP Detectors

In the original MVSO image mining,¹ when we downloaded data, we first queried the Flickr API using our multilingual ANPs. Generally, there are two types of ways to search for photos on Flickr: asking the Flickr API to search for the query in the (1) tags of a photo, or (2) anywhere in the photo text data, including photo title, description and tags. We call these “tag search” and “free-text search,” respectively. Tag search yields less but more precise results, and free-text search will give more but noisier data, i.e. leading to weaker supervision. The intuition is that if a user uses an ANP as a tag, then it is more likely to describe what is actually visually present in the image, while an ANP that occurs in, say the description, may be relevant, but not visually present. In our data collection, we used an upper bound of 1,000 images per ANP where we first queried via tag search and if the upper bound had not been reached, we then pulled results from free-text search on Flickr. As a result, the original MVSO image dataset constitutes a *hybrid-pool* of Flickr images with respect to how the querying was performed. The six ANP detector banks presented in the original release¹ were based on this very hybrid-pool of images. In Table 1, we show the top- k accuracies of the original¹ CaffeNet-structured,¹¹ DeepSentiBank-fine-tuned³ detector banks using this hybrid-pool of images. Likewise, in Table 2, we show the top- k performances of the new Inception-styled architectures for the same six languages on the same hybrid-pool of images. Top- k accuracy refers to the percentage of classifications for which the true class is in the top k predicted ranks.

Language	#ANPs	#params	#train	#test	time	top-1	top-5
English	4,342	30.50	3,236,728	807,447	370	13.64%	26.63%
Spanish	2,382	20.84	1,085,678	270,400	106	13.86%	27.98%
Italian	1,561	7.57	602,424	149,901	57	17.46%	32.35%
French	1,115	7.12	462,522	115,112	36	16.76%	34.27%
German	275	6.26	108,744	27,048	11	31.08%	54.10%
Chinese	243	6.22	102,740	25,575	10	25.96%	47.11%

Table 2: Classification accuracies of the new ANP detector banks (%) using an Inception-based architecture⁴ with training walltimes (hours) and number of parameters (millions). The classes, training and testing sets all match those from Table 1, i.e. those used in the original MVSO release.¹

Compared to the previous detector banks in Table 1 which we fine-tuned from an affectively biased set of network weights,³ training takes noticeably longer in most cases since we now fine-tune from a model trained on ILSVRC. Nonetheless, we observe that we are able to get improved ANP classification rates on most languages with Inception,⁴ particularly when there is more data for the networks to take advantage of; for example, with English ANP detection, we get a 35.05% relative improvement at top-1. Overall, except for French, we see a consistent improvement at top-5, indicating that more semantically relevant ANPs are being surfaced into the top ranks by Inception compared to the CaffeNet model. While we do not see this same consistency at top-1, given that these detectors currently treat all ANPs as if they came from a flat taxonomy, many of the ANPs may in fact be semantically close to each other as indicated by the hierarchical grouping studied in the original ontology construction process.¹ As a result, from a purely empirical standpoint, a top-1 accuracy metric may treat a ‘pretty flower’ prediction for a ‘beautiful flower’-labeled image to be a misclassification. A top-5 metric in these ontology-structure-agnostic settings then may be a more reasonable indicator of detector bank performance than a top-1 performance metric.

4.2 Tag-pool Multilingual ANP Detectors

Since the image corpus in the hybrid-pool setting, where images come from both tag and free-text search on Flickr, can be noisy, we perform a separate set of experiments where we train our ANP detector banks just from the subset of images that come from tag-restricted queries on Flickr. We call this the *tag-pool* of MVSO images. Though it is a subset, each ANP still has a maximum of 1,000 images to learn from and we split the dataset as before 80/20% per ANP. However, since we still enforce the requirement that there are at least 125 images per ANP, some languages experienced significant decrease in ANP coverage. In Figure 2, we show the change in ANP and image coverage when going to a tag-restricted subset in MVSO. We observe that English experiences a small 9.92% loss in ANPs and a 63.04% loss in image count when restricting to tag-only queries while all other languages experience about 80% coverage loss for both ANPs and images. Though the ANPs in




Figure 2: Hybrid vs. Tag-pool MVSO Coverage. In (a), we show the coverage of ANPs and in (b), images, across the six major languages in MVSO.¹ English suffers the least from the tag restriction while all other language experience a large drop in ANP and image coverage.

MVSO are useful given their pervasive occurrence in social media, this indicates that there are a fair number of images in hybrid-pool have a considerable amount of weak supervision. At the same time, it is worth noting that it is wrong to assume all images from the free-text search are not useful for visual recognition of adjective-noun pairs.

In Table 3 below, we show the detector bank performances for the six major languages using the tag-restricted subset of images with CaffeNet-styled¹¹ architectures. To train these tag-pool ANP detectors, we follow training scheme from the hybrid-pool ANP detectors in the original MVSO release¹ for 1-to-1 comparison, pre-training with DeepSentiBank³ using Caffe¹¹ with the same optimization strategies, with the only slight difference being that we use a different GPU, i.e. 980 versus Titan X, and use the true training mean image instead of the ILSVRC mean. Again, note that the number of ANP classes each detector learns over has reduced compared to those seen earlier in Table 1 due to the tag-based pool restriction.

Language	#ANPs	#params	#train	#test	time	top-1	top-5
English	3,911	72.89	2,294,411	255,192	85	19.00%	33.81%
Spanish	502	58.92	164,119	18,276	12	29.07%	52.86%
Italian	402	58.52	151,088	16,807	9	33.69%	55.76%
French	222	57.78	80,790	8,986	6	34.70%	63.32%
German	90	57.24	32,195	3,578	2.5	47.04%	74.62%
Chinese	60	57.11	19,044	2,122	1.5	45.05%	71.35%

Table 3: Classification performances of the CaffeNet-based¹¹ ANP detectors (%) using images retrieved by restricting Flickr API queries to tag-based search. Note that walltimes (hours) are for running a fixed 80 epochs.

At first glance, the top-1 and top-5 accuracy rates seem much higher than those we observed with a similar architecture in Table 1. However, it is still difficult to determine if restricting to the tag-pool actually contributed to more reliable labels and thus better performing ANP detectors, especially since the number of output classes has changed. As a result, we evaluated the hybrid-pool ANP detectors from §4.1 on the test set of the tag-restricted pool of images as well as the the intersection of the test sets. Since random shuffling was done during train/test splitting for both hybrid- and tag-pool image datasets, the latter ensures that no images in the test set (#test) were used in either model’s training. The results from this evaluation are shown in Table 4. We observe that in both cases, the classification performances are both lower than those in Table 3 by about 10% absolute on each language, indicating that the tag-pool ANP detectors do indeed achieve a comparably higher accuracy on the same set of ANPs. In addition, for an even closer comparison, we tested the trained tag-pool ANP detectors on the intersection of the hybrid-pool and test-pool test sets from Tables 1 and 3, respectively, as shown in Table 5. The results here are thus most comparable to those found in Table 4(b). Although there are significantly less test images than with Table 1, 3, or 4(a), we observe that the top-1 classification rates for the tag-pool English ANP detector bank is a relative 59.21% improvement over the hybrid-pool English ANP

Language	#ANPs	#test	top-1	top-5	#test	top-1	top-5
English	3,911	756,243	10.47%	21.67%	51,062	12.04%	24.58%
Spanish	502	91,002	14.86%	30.48%	3,740	17.91%	34.39%
Italian	402	64,668	22.51%	39.44%	3,386	24.25%	42.85%
French	222	38,107	21.78%	43.10%	1,835	25.29%	46.54%
German	90	13,249	36.18%	58.99%	769	39.14%	62.68%
Chinese	60	9,093	29.47%	48.78%	412	34.95%	51.94%

(a)

(b)

Table 4: Classification performances of the original CaffeNet-based¹¹ ANP detectors (%) from the ACM Multimedia 2015 release,¹ i.e. same detectors as in Table 1, but now evaluated only on the test images **(a)** of the tag-pool in Table 3, and **(b)** intersection of both the hybrid-pool and tag-pool test sets.

detector bank on a comparable dataset. ANP detection in all other language likewise see an improvement in the evaluation using the tag-pool ANP detectors.

From the evaluations in Tables 3-5, we believe we can conclude that training ANP detector banks using the tag-pool have led to significantly better ANP detectors. The trade-off here, however, is that we have dramatically less images to train on and also a significantly reduced coverage of ANPs per language, except in the case of English. Nonetheless, all this serves as concrete evidence that tag-based labels are more reliable than those mined from the metadata overall since users are more likely to use tags to describe content that is visually present in the image. While this does not mean that all labeled images gathered from free-text metadata context are wrong, it does reduce the trust that should be placed on their ability to represent elements visually present in the associated image. We believe that in the future development of ANP detectors factoring such label reliability will allow for a greater coverage of ANPs without sacrificing in the area of data sparsity or the classification performance of ANP detector banks.

Language	#ANPs	#test	top-1	top-5
English	3,911	51,062	19.17%	34.01%
Spanish	502	3,740	28.85%	52.75%
Italian	402	3,386	32.75%	54.90%
French	222	1,835	35.80%	64.20%
German	90	769	49.41%	74.51%
Chinese	60	412	41.50%	67.23%

Table 5: Classification performances of the CaffeNet-based¹¹ ANP detectors (%) using tag-restricted images but now evaluated only on the set of test images intersecting the hybrid pool of Table 1, i.e. original MVSO data release,¹ and tag-restricted pool of Table 3. These tag-restricted performances are thus most comparable to the results reported in Table 4(b).

5. CONCLUSION

In this technical report, we have presented new ANP detector banks for MVSO for six major languages using a more modern network architecture than in the original MVSO release. These detectors have a reduced model memory footprint compared to their predecessors and achieve very competitive or even significantly better ANP detection performance, e.g. a 35.05% improvement for English ANP detection. In addition, we presented ANP detectors that use a restricted subset of images and ANPs that originate only from tag-based Flickr API retrieval results, which resulted in much higher ANP classification performances due to improved label quality. We also publicly release the new Inception-based ANP detector banks, tag-pool CaffeNet ANP detector banks, and the list of image URLs corresponding to the tag-restricted MVSO image subset used in this report.

In the future, we would like to investigate additional fine-tuning strategies for training ANP detectors. For example, in the original release,¹ we noted that fine-tuning from DeepSentiBank³ was done in the hopes that we initialized the models from an affectively biased source; however, it is not empirically clear whether such a biasing actually helps or hurts detector performance. In addition, we would like to investigate and benchmark

our ANP detectors against other popular network architectures. Also, we plan to investigate how we might integrate the semantic hierarchy in MVSO to train ANP detectors that exploit the tree-structure.

ACKNOWLEDGMENTS

We would like to thank Miriam Redi, Mercan Topkara, Nikolaos Pappas and Tao Chen from Multilingual Visual Sentiment Ontology (MVSO) team for their continued support and insightful discussions. Also, we thank Margaret Yuying Qian for performing some of the early tag-restricted subset partitioning and statistics.

REFERENCES

- [1] Jou, B., Chen, T., Pappas, N., Redi, M., Topkara, M., and Chang, S.-F., “Visual affect around the world: A large-scale multilingual visual sentiment ontology,” in [*ACM Multimedia (MM)*], (2015).
- [2] Krizhevsky, A., Sutskever, I., and Hinton, G. E., “ImageNet classification with deep convolutional neural networks,” in [*Neural Information Processing Systems (NIPS)*], (2012).
- [3] Chen, T., Borth, D., Darrell, T., and Chang, S.-F., “DeepSentiBank: Visual sentiment concept classification with deep convolutional neural networks,” *arXiv preprint arXiv:1410.8586* (2014).
- [4] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A., “Going deeper with convolutions,” in [*Computer Vision and Pattern Recognition (CVPR)*], (2015).
- [5] Picard, R. W., [*Affective Computing*], MIT Press (1997).
- [6] Borth, D., Ji, R., Chen, T., Breuel, T., and Chang, S.-F., “Large-scale visual sentiment ontology and detectors using adjective noun pairs,” in [*ACM Multimedia (MM)*], (2013).
- [7] Liu, H., Jou, B., Chen, T., Topkara, M., Pappas, N., Redi, M., and Chang, S.-F., “Complura: Exploring and leveraging a large-scale multilingual visual sentiment ontology,” in [*Intl Conf. on Multimedia Retrieval (ICMR)*], (2016).
- [8] Pappas, N., Redi, M., Topkara, M., Jou, B., Liu, H., Chen, T., and Chang, S.-F., “Multilingual visual sentiment concept matching,” in [*Intl Conf. on Multimedia Retrieval (ICMR)*], (2016).
- [9] Plutchik, R., [*Emotion: A Psychoevolutionary Synthesis*], Harper & Row (1980).
- [10] Jou, B., Qian, M. Y., and Chang, S.-F., “SentiCart: Cartography & geo-contextualization for multilingual visual sentiment,” in [*ACM International Conference on Multimedia Retrieval (ICMR)*], (2016).
- [11] Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T., “Caffe: Convolutional architecture for fast feature embedding,” in [*ACM Multimedia (MM)*], (2014).
- [12] Borth, D., Chen, T., Ji, R., and Chang, S.-F., “SentiBank: Large-scale ontology and classifiers for detecting sentiment and emotions in visual content,” in [*ACM Multimedia (MM)*], (2013).
- [13] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L., “ImageNet Large Scale Visual Recognition Challenge,” *Intl Journ. of Computer Vision (IJCV)* **115**(3) (2015).
- [14] Lin, M., Chen, Q., and Yan, S., “Network in network,” *arXiv preprint arXiv:1312.4400* (2014).